

Information-Optimal Selective Data Return for Autonomous Rover Traverse Science and Survey

David R. Thompson, Trey Smith and David Wettergreen

Abstract—*Selective data return* leverages onboard data analysis to allocate limited bandwidth resources during remote exploration. Here we present an adaptive method to subsample image sequences for downlink. We treat selective data return as a compression problem in which the explorer agent transmits the subset of measurements that are most informative with respect to the complete dataset. Experiments demonstrate selective downlink of navigation imagery by a rover during autonomous geologic investigations in the Atacama desert of Chile. Here automatic analysis identifies informative images using classifications based on natural image statistics. Image texture analysis, together with a context-sensitive Hidden Markov Model representation, permits adaptive downlink in response to geologic unit boundaries. Selective data return improves the science content of returned data for this geologic mapping task.

I. INTRODUCTION

Autonomous robots performing remote science exploration tasks can improve their performance by adapting to the content of collected data. Autonomous onboard data analysis will become increasingly important for rovers traveling long distances between command cycles. Multi-kilometer traverses will collect data from large areas of terrain that will never be seen by humans [19]; these rovers can prioritize the most important data products for their limited-bandwidth communication with Earth [2].

The formal objectives of selective data return will depend on science mission requirements. Scientists seeking evidence of specific phenomena can use a *target signatures* approach [2], commanding the explorer agent to prioritize these features for downlink. For example, the E0-1 orbiter uses intelligent data return that targets transient events such as floods and volcanoes [5]. Image analysis onboard the Mars Exploration Rovers identifies dynamic atmospheric phenomena, such as clouds and dust devils, and then flags these images for selective return [3]. Researchers have developed more sophisticated joint utility functions that combine multiple targets of interest [17] and predict scientists' preferences over sets of returned features [8].

Target signature strategies may not be appropriate for all applications. Many science features are too subtle for real-time pattern recognition to identify consistently. Moreover, features' science value is often sensitive to the local spatial context; the same feature can be more or less interesting



Fig. 1. The rover platform “Zoë” used in Atacama desert autonomous geology experiments. Navigation cameras used for capturing image sequences are mounted midway along the rover mast.

depending on the location in which it is found. These distinctions are especially important in site survey applications that aim to characterize large areas efficiently.

Here we explore an alternative data return strategy that addresses these challenges. We formulate selective data return as a compression problem with the goal of transmitting as faithfully as possible the entire content of collected data. Our objective function minimizes scientists' posterior uncertainty about the complete set of measurements given the returned data. The Shannon entropy links our approach to the tradition of optimal compression and communications theory [7], as well as adaptive sampling strategies such as active learning [15] and Bayesian experimental design [4].

This paper addresses the specific problem of selective data return in an ordered image sequence. Image sequences are common in planetary exploration applications — they occur in flybys, landing sequences, and rover traverse datasets. These image sequences can usually be collected at a much higher rate than return bandwidth can accommodate. However, visual object detection is a challenging pattern recognition problem so it is usually difficult to characterize their science content accurately. In our approach a Hidden Markov Model (HMM) compensates for this uncertainty by leveraging the local continuity of the explored environment. We demonstrate selective downlink of rover navigation imagery from autonomous geologic investigations in the Atacama desert of Chile. The rover uses image texture analysis together with a context-sensitive HMM representation in order to identify and communicate the location of geologic units.

This work supported by NASA grants NNG0-4GB66G and NAG5-12890
D. R. Thompson and D. Wettergreen are with the Robotics Institute, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213, USA. {drt,dsw}@ri.cmu.edu
T. Smith is with Carnegie Mellon West / NASA Ames Research Center, Mail Stop 269-3, Moffett Field, CA 94035-1000, USA. {trey.smith}@west.cmu.edu

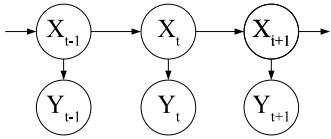


Fig. 2. Hidden Markov Model of rover traverse imagery. Image content is represented by hidden variables $\{X_1, X_2, \dots, X_T\}$ associated with proxy features $\{Y_1, Y_2, \dots, Y_T\}$.

II. APPROACH

The remote exploration procedure consists of two stages. In the initial *data collection* phase the explorer agent collects a series of images $A = \{M_1, M_2, \dots, M_T\}$. Each image M_t is associated with the independent experimental variable t . We will consider the one-dimensional case where t is the observation's location in an ordered series indexed by time or position along a transect.

The images vary across every pixel value so each M_t is drawn from a very high dimensional space. However, most of this information is superfluous to the science investigation. Instead, remote scientists base their interpretations on more abstract features such as physical objects in the scene. For example, a remote geologic site survey might involve differences in sediment structure [1], shape or size attributes of rocks [13], or the presence of desired objects like outcrops. We represent these attributes by assigning each image a discrete-valued label X_t so that the dataset maps to a sequence of variables $B = \{X_1, \dots, X_T\}$. These values represent the “high-level content” that would be readily apparent to a scientist interpreting the images.

The exploration agent attempts to identify this science content with automatic pattern recognition. These proxy classifications Y_t yield a sequence $C = \{Y_1, \dots, Y_T\}$. Since the automatic classifications are imperfect the remote explorer agent cannot be certain of the image's true label. However, the agent can still exploit noisy correlations with science phenomena to make principled data return decisions.

In the *data return* phase the agent computes a downlink transmission B' to optimize a reward function $R(B'; B)$. This reward should reflect the downlink's science value. We formulate subset selection as a compression task to transmit B to a receiver with minimal distortion across a bandwidth-limited channel [7]. In other words, the agent should favor downlinks that preserve the science content of the complete image set. We can write the reward as a functional over posterior distributions: $R(P(B|B'))$. We desire that the reward functional encourage posterior certainty about B ; it should be some measure of information gain such as the negative conditional Shannon entropy $-H(B|B')$. This work employs the negative sum of residual marginal entropies $\sum_t H(X_t|B')$. We will see that this reward functional facilitates an easy solution to the optimization problem.

While in general we are interested in an arbitrary transmission that maximizes $R(P(B|B'))$, circumstances of remote science exploration favor a *subset selection* approach to image return. Remote science missions often place a high

value on lossless data products that can serve as validated, archival science records. Here we will follow this convention and return a subset of collected measurements. For a cost \mathcal{L} reflecting the size of the transmission, the penalty of sending the classification results will be trivial relative to the cost of the images themselves: $\mathcal{L}(M_t) \gg \mathcal{L}(Y_t)$. This suggests that image classifications will be returned in any case, and that the loss function can be expressed in the form $R(P(B|B', C))$. For an available downlink budget k the objective becomes:

$$\begin{aligned} B'_{\text{opt}} &= \sup_{B'} R(P(B|B', C)) \\ B' &\subseteq B \\ \mathcal{L}(B') + \mathcal{L}(C) &\leq k \end{aligned} \quad (1)$$

Choosing the optimal return subset requires modelling correlations between Y_t and the actual label X_t . However, one can also exploit relationships between hidden variables themselves. One important relationship concerns the spatial continuity of the explored environment — consecutive images are likely to have the same label. We capture these correlations using a Hidden Markov Model (HMM) where each image in the sequence is conditionally independent of the others given its immediate neighbors (Fig. 2). Model parameters consist of initial state probabilities $P(X_1)$, transition probabilities $P(X_{t+1}|X_t)$ representing the probability of state transitions from one timestep to the next, and emission probabilities $P(Y_t|X_t)$ representing correlations between images' true labels and the automatic classifications.

Designers have several options for estimating HMM parameters. If scientists can enumerate all the image features of interest then HMM parameters can be estimated from labeled training data using standard supervised learning techniques. The Maximum Likelihood parameter estimates for discrete labels are given by the empirical counts:

$$\begin{aligned} \hat{P}(X_{t+1} = i | X_t = j) &= \frac{\sum_t \delta(X_{t+1} = i, X_t = j)}{\sum_t \delta(X_t = j)} \\ \hat{P}(Y_t = k | X_t = j) &= \frac{\sum_t \delta(Y_t = k, X_t = j)}{\sum_t \delta(X_t = j)} \end{aligned} \quad (2)$$

In the following experiments we will assume that previous training data is available so that the HMM parameters are known in advance.

On the other hand, if the investigation is a general site survey the scientists might not know what to expect and the phenomena of interest represented by X_t may be ambiguous. In this case it might be better to estimate labels automatically at runtime. For example, the explorer agent could use the Baum-Welch algorithm [16] to estimate Maximum Likelihood HMM parameters. The best number of state labels is a model selection problem to be addressed by cross validation or measures such as the Akaike Information Criterion [18].

Following Krause and Guestrin [14] we base our reward function on the sum of posterior marginal entropies. The complete reward decomposes into the sum of individual rewards for each X_t .

$$\begin{aligned}
R(P(B|B', C)) &= \sum_{X_t \in B} R_i(P(X_t|B', C)) \\
&= - \sum_{X_t \in B} H(X_t|B', C)
\end{aligned} \quad (3)$$

Expected conditional entropies of the hidden variables can be calculated in the following manner, where x represents the possible assignments to variables in B and x' represents assignments to B' .

$$\begin{aligned}
R(P(B|B', C)) &= - E_{B'} \sum_{X_t \in B} H(X_t|B', C) \\
&= - \sum_{x'} P(x') \sum_{x_t, x'} P(x_t, x'|C) \log P(x_t|x', C)
\end{aligned} \quad (4)$$

Krause and Guestrin present a dynamic programming algorithm for optimal subset selection that exploits a decomposition of the local reward function [14]. We omit the details here for brevity; the basic idea is to decompose the reward into subchains that become conditionally independent when hidden nodes are revealed. Ultimately we need only evaluate combinations of possible assignments to two variables in B' - those at either end of the current subchain. The resulting recursive algorithm can be used to compute $R(P(B|B', C))$.

In summary, selective return consists of the following procedure:

- 1) The agent performs a traverse and collects images whose science content is a sequence of hidden variables $B = \{X_1, X_2, \dots, X_T\}$.
- 2) Onboard pattern recognition processes the images to yield the sequence of noisy classifications $C = \{Y_1, Y_2, \dots, Y_T\}$.
- 3) The agent chooses the subset of images B' to optimize $R(P(B|B', C))$ from equation 4.
- 4) Scientists receive the image subset and observe the values $X_t \in B'$.

A series of simulated trials investigated the effects of different environments and classification noise on data return behavior. We synthesized data from a traverse sequence in which $X_t \in \{0, 1\}$ and $Y_t \in \{0, 1\}$. Each X_t maps onto its corresponding binary-valued classification with a small probability of error. Similarly, the environment has some small probability of a state transition between sampling locations.

Fig. 3 illustrates typical data return behavior with a simulated traverse containing a single state transition. The agent mitigates observation uncertainty by returning images surrounding the likely transition; this helps the receiver to identify its precise location. Fig. 4 shows a traverse with equivalent states and observations. In this second case the agent's model parameters imply frequent state transitions and a poor correlation between science content and the automatic classification. The agent considers its proxy classifications to be less informative and opts for a more conservative policy. Its optimal strategy is closer to periodic sampling with

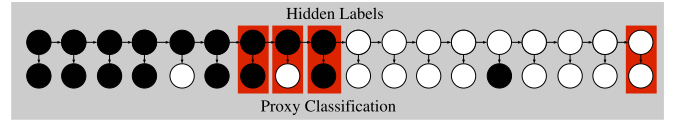


Fig. 3. Typical behavior of the data return algorithm in the noisy binary case, with nodes colored according to their value. The agent only observes the proxy classifications Y_i on the bottom row. Rectangles indicate images selected for return. This example uses transition and error probabilities of 0.1, for which the optimal subsampling consists of images clustered near likely state transitions.

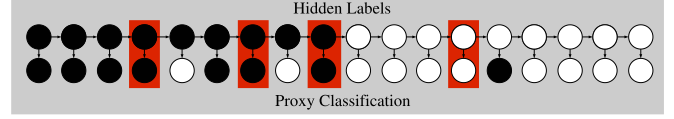


Fig. 4. In this segment the agent presumes transition and error probabilities of 0.33, but the observations are identical to Fig 3 above. Increased noise implies that automatic classifications are less informative, resulting in a more conservative subsampling strategy.

returned images spaced regularly throughout the sampling domain.

We computed optimal data return strategies for a series of simulated traverses in which 30 collected images were subsampled to a 6-image downlink dataset. Fig. 5 characterizes the resulting downlinks for various levels of transition and observation noise. Isocontours show the variance in the distance between samples; high variance suggests irregular clusters of samples such as those in Fig. 3 while low variance indicates evenly-spaced sampling as in Fig. 4. Formulating selective data return as compression prescribes a return strategy that reflects both the context of local observations and the fidelity of onboard pattern recognition to the underlying phenomena of interest.

III. FIELD EXPERIMENTS

This section details our implementation of the selective data return algorithm for geologic boundary detection. We

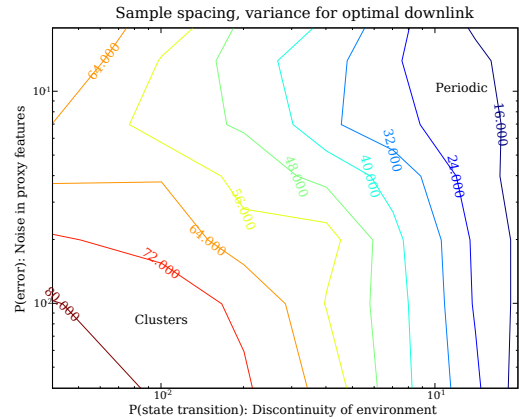


Fig. 5. Variance in spacing of optimal image return as a function of noise in the environment and onboard classification.

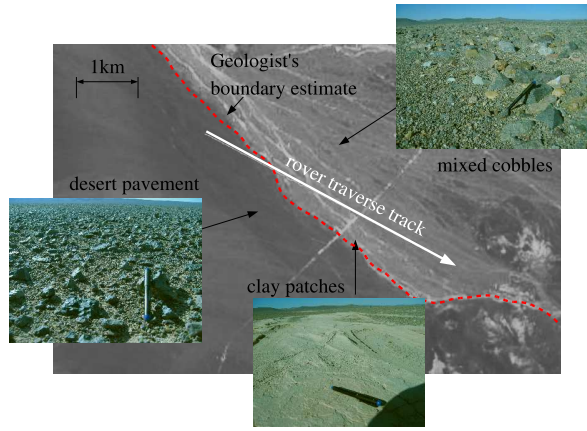


Fig. 6. ASTER orbital imagery corresponding to a 4.3km traverse. The initial terrain consists of homogeneous desert pavement covered with small rocks (< 10 cm). After the first kilometer the terrain changes (yellow dashed line) to a mixture of larger block sizes with occasional bare clay. Scientists evaluating the orbital data hypothesized a second boundary (red dashed line) but no change was clearly visible from the ground. In this annotated image a thin white line shows the traverse locations where navigation imagery was captured.

consider a series of autonomous rover traverse experiments in the Atacama Desert, a Mars-analog site in Chile. These experiments deal with operations at the “Guanaco” site (referenced elsewhere as Site F) that involved over 50km of autonomous traverse. The rover platform used for the experiments is “Zoë,” an exploration robot developed at Carnegie Mellon [19]. Zoë is a capable vehicle that travels up to 1 m/s in open terrain (Fig. 1). Onboard stereo vision, obstacle avoidance and path planning enable autonomous traverses of up to several kilometers per day.

During field operations scientists studying Atacama geology guide the rover remotely by selecting a series of waypoints in satellite images. Occasionally these traverses cross an area hypothesized to contain a contact between two or more geologic units. Fig. 6 shows one unit boundary where a multi-kilometer drive action crosses a clear border between a low-albedo region of desert pavement and a high-albedo region of mixed cobbles.

These surface unit borders are geologically informative but difficult to verify from orbital imagery alone. However, in planetary exploration scenarios it would be infeasible to identify unit boundaries by transmitting the entire catalog of collected images. Here we address the issue with adaptive data return. The HMM provides a natural framework for mapping geologic state during long-distance autonomous traverse. We take the science content X_i to be a label of the geologic unit for image M_t .

A. Image Processing

The navigation cameras used for this experiment provide a 60-degree field of view at 320x240 pixel resolution; they point downward to view the terrain directly in front of the rover. The image analysis scheme classifies these images according to the geologic unit their contents represent. Image texture has been found in the laboratory to be a good

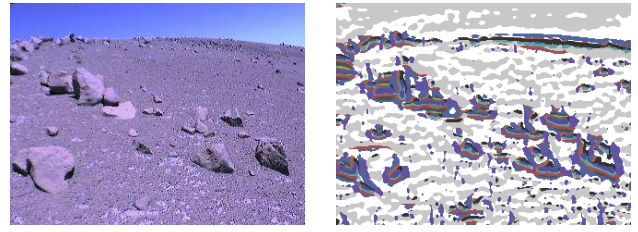


Fig. 7. Original navigation image (left) and false color image showing texton values of individual pixels (right). To limit horizon and sky artifacts, we use only the bottom half of each image to compute its summary descriptor.

correlate of geomorphology [9]; our analysis uses simple local texture descriptors extracted from the foreground of each image.

We describe local textures using the texton method of Varma and Zisserman [21]. We convolve all images in a training set with the Maximum Response 8 filter bank. This results in an 8-dimensional response vector for each pixel. The responses from several training images are aggregated and clustered using k-means to form a set of 16 universal textons. These textons yield an image map that assigns each individual pixel to its Euclidean-nearest texton (Fig. 7). Different textons approximately capture the different kinds of surface materials such as rocks, shadows, sediments of varying texture, and salt deposits.

The proportions of each texton form a 16-dimensional feature vector which serves as a feature space for predicting the images’ geologic unit type. We produce training data automatically by indexing rover position against a georeferenced map of predicted unit boundaries drawn previously by geologists using orbital images. The rover’s uncertainty about image labels is negligible whenever the rover’s distance from all boundaries is large relative to its own localization uncertainty. In cases where geologists are fairly certain about the boundary position we can assign an unambiguous unit label to those images. This training data can be reserved for future traverses where the boundary locations are less certain. Image classification uses a support vector machine with a radial basis kernel function [6].

At the Guanaco site expedition, four extended traverses

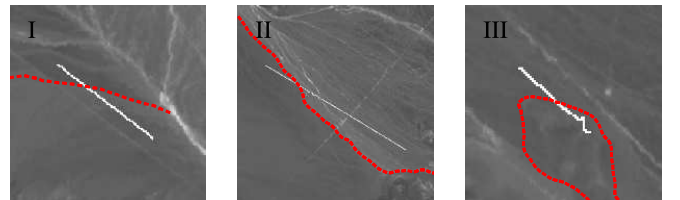


Fig. 8. Three boundary-crossing traverses from the Atacama expedition. These Visible/Near Infrared orbital images were collected by the ASTER instrument [22] and used by geologists to draw preliminary geologic maps. The bright white line in each image represents the position track of the rover during the traverse, while the red dashed line shows the boundary hypothesized by the project geologist. Each image pixel represents approximately 15 square meters of terrain. Desert roads, drainages, and other high-albedo features are also visible.

TABLE I
CLASSIFICATION ACCURACY FOR TEXTURE-BASED IMAGE FEATURES.

Traverse	Training Accuracy	Test Accuracy	State Estimation
I	96.8	82.0	84.0
II	100.0	88.5	96.2
III	97.4	88.3	98.1

cross boundaries that remote geologists have previously identified as probable contact points between units of surface material. We use one of these traverses as a test set for tuning image processing scheme and apply our selective return strategy to the remaining three traverses (Fig. 8).

Table I shows the results of the classification procedure. The first column shows 10-fold cross validation training error, while the second shows test accuracy on over 1000 images from the Atacama image dataset. Test accuracy is substantially lower than cross-validation would suggest due to several factors that complicate generalization across traverses. The geologic units themselves are not completely homogeneous, and the geologic content of two traverses may differ despite having the same map label. Moreover, changing lighting conditions create shadows and cause systematic changes in the visual appearance of terrain. Two traverses may therefore exhibit markedly different features if they occur at a different time of day. However, for the purposes of this experiment it is not necessary that the image classification be perfect; rather, our goal is to account for this error to improve selective downlink.

The final column of table I shows the image label accuracy on test results after posterior smoothing of the HMM. The Viterbi algorithm [10] produces the Maximum Likelihood state sequence, which is a better correlate of the true unit labels than the independent image classifications. We include this score for interest only, to suggest that the HMM structure helps state estimation by capturing the continuity of the real environment. The Viterbi parse is an estimate of X_t and not Y_t ; it does not necessarily reflect the accuracy of the image classification procedure itself.

B. Selective Return

Each of the three test traverses consist of a navigation sequence of 50 images; selective return produces a subset of 10 images for downlink. We use conservative HMM parameters: a transition probability of $P(X_{t+1} = i | X_t = i) = 0.9$, and a classification accuracy of 0.8. The transition probability is fixed through trial and error to produce reasonable results for the training traverse. Note that this value is sensitive to the imaging rate; a denser image sequence would imply a lower state transition probability. Figure 9 shows the resulting locations of the returned images, the location of the boundary hypothesized by geologists, and the rover's posterior probability estimate of the hidden state.

Of the three segments, the state estimation in traverse I exhibits the worst match to the geologist-hypothesized boundary. We attribute this to particularly poor classification scores (Table I). The distinctions between the geologic units

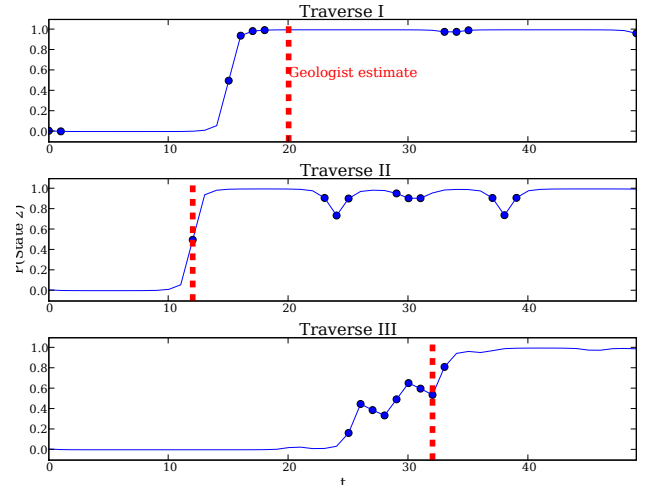


Fig. 9. Returned images from the three traverses, along with the posterior state estimate. Selective return allocates images to areas where the hidden state is most ambiguous, such as border areas (top and bottom) and anomalous observations (middle)

was subtle, and the units themselves were highly heterogeneous (each containing mixtures of small of rock-strewn and open areas). Thus, the actual location of the true boundary is somewhat ambiguous. Moreover, field conditions complicated the selective return procedure. The traverse occurred late in the day with rapidly changing light conditions that independently altered image appearance over time.

Traverse II's state estimation accurately identifies the geologic unit boundary. The return procedure fixes any remaining ambiguity with an image that corresponds exactly to the anticipated border. Other images selected for return include those from three locations later in the traverse. Here the system classifies several images as belonging to the first unit, which increases the probability of a rare state transition. The resulting uncertainty about hidden state in these areas motivates the selective return procedure to spend extra bandwidth there.

In traverse III geologists have hypothesized a homogeneous traverse finished by a single discrete transition at the base of a hill. In fact, the approach to the hill was itself strewn with rocks similar to those present at the peak. This is reflected by a gradual transition rather than a clear boundary. It is unsurprising that image labels in this area are mixed, resulting in an uncertain, gradually changing state estimate. The selective return procedure reacts by allocating samples more-or-less evenly throughout this transition area. A clear difference in sedimentology is evident between the first and last images from the downlinked dataset (Fig. 10).

The field experiments demonstrate that time sequence models are appropriate for remote transect surveys. The HMM permits context-sensitive state inference and calculation of informative observations for selective data return.

IV. CONCLUSIONS

This paper presents a new approach for selective return of ordered data products during remote exploration. We treat

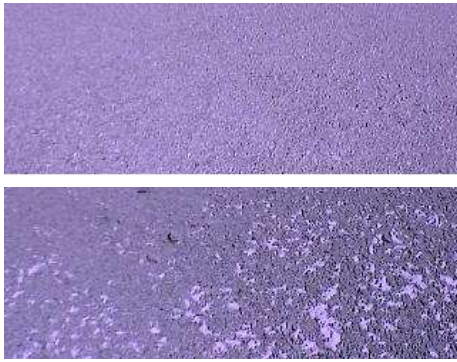


Fig. 10. The first (top) and last (bottom) images returned from traverse III. Only the analyzed foreground portion of each image is shown.

the problem as a compression task where the agent has limited information about the science content of collected data. In these cases, noisy feature classification can still be used to inform data return decisions. The performance of this classification is an essential consideration, since accurate feature detection permits highly selective targeted sampling. As feature classifications become less informative, conditional dependencies due to local context begin to dominate the observations. The agent's behavior comes to resemble the *a priori* data return strategy (in our case, periodic sampling).

There are multiple avenues for improving the demonstrated system. With respect to geologic investigations, better image analysis is a promising means to increase performance. We aim to leverage recent work in the computer vision community on scene classification and natural image statistics. Better patch descriptors and hierarchical content models are likely possibilities. It may be desirable to introduce explicit detection of scene objects such as rocks, sediment patches, or layered outcrops. Finally, the classification difficulties we experienced with traverse I suggest it will be necessary to address generalization issues such as shadow removal, color constancy and lighting invariance. The model itself is another candidate for future work. Alternative state representations, such as continuous states, may be more appropriate for some domains. Since the inference algorithms we employ here apply to any chain-structured model, distributed representations such as Factorial Hidden Markov Models [12] might permit more complicated and structured science content. Gaussian processes [20] are also promising possibilities for time series data.

Several open challenges remain. We have demonstrated that classification inaccuracy can have a significant impact on sample spacing. This raises the question of how to estimate future classification performance. By definition, exploration tasks involve investigating new areas and phenomena that differ from our expectations. How then shall we accurately predict generalization error of the onboard pattern recognition? It may be possible for the agent itself to recognize anomalous situations by tracking the data likelihood [11]. If the traverse likelihood lies within the envelope of its training data, the agent can be more confident that its targeted

sampling is improving performance.

V. ACKNOWLEDGMENTS

The authors thank the Life in the Atacama project team including Dominic Jonak and Michael Wagner, and the science team for their assistance with geologic interpretations. NASA ASTEP grant NNG0-4GB66G supported the work.

REFERENCES

- [1] N. A. Cabrol, J. D. Farmer, E. A. Grin, L. Richter, L. Soderblom, R. Li, K. Herkenhoff, G. A. Landis, and R. E. Arvidson, "Aqueous Processes at Gusev Crater Inferred From Physical Properties of Rocks and Soils along the Spirit Traverse," *Journal of Geophysical Research* 111, E02S20, 2006.
- [2] R. Castaño, R. C. Anderson, T. Estlin, D. DeCoste, F. Fisher, D. Gaines, D. Mazzoni, and M. Judd, "Rover Traverse Science for Increased Mission Science Return," *Proceedings of the IEEE Aerospace Conf.*, Big Sky, Montana, 2003.
- [3] A. Castaño, A. Fukunaga, J. Biesiadecki, L. Neakrase, P. Whelley, R. Greeley, M. Lemmon, R. Castano, and Steve Chien. "Automatic detection of dust devils and clouds on Mars." *Machine Vision and Applications, Special Issue*, 2007.
- [4] K. Chaloner and I. Verdinelli, "Bayesian Experimental Design: A Review," *Statistical Science* 10:3, 1995.
- [5] Steve Chien et al. "The EO-1 Autonomous Science Agent," *Conference on Autonomous Agents and Multi-Agent Systems*. New York City, USA. July 2004.
- [6] C. Chang and C. Lin, *LIBSVM: a library for support vector machines*, 2001. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>
- [7] T. M. Cover and Joy A. Thomas. *Elements of Information Theory*. New York: Wiley and Sons, 1991.
- [8] M. desJardins, E. Eaton, and K. L. Wagstaff. "Learning User Preferences for Sets of Objects," *Proceedings of the Twenty-Third International Conference on Machine Learning*, June 2006.
- [9] H. Dunlop, *Automatic Rock Detection and Classification in Natural Scenes*, Masters Thesis, Carnegie Mellon University, CMU-RI-TR-06-40, August 2006.
- [10] G. D. Forney. "The Viterbi Algorithm," *Proceedings of the IEEE* 61, 1973. p. 268–278.
- [11] M. Fox, J. Gough, and D. Long. "Detecting Execution Failures Using Learned Action Models," *AAAI Proceedings*, 2007.
- [12] Z. Ghahramani and M. I. Jordan, "Factorial Hidden Markov Models," *Advances in Neural Information Processing Systems* 1995.
- [13] M. B. Golombek, et. al. "Geology of the Gusev Cratered Plains from the Spirit Rover Transverse," *Journal of Geophysical Research*, 111, E02S07, 2006.
- [14] A. Krause and C. Guestrin. "Optimal Value of Information in Graphical Models - Efficient Algorithms and Theoretical Limits," *Proc. of the International Joint Conference on Artificial Intelligence*, 2005.
- [15] D. MacKay. "Information-Based Objective Functions for Active Data Selection," *Neural Computation*, 4:4, 1992.
- [16] R. L. Rabiner. "A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition," *Proceedings of the IEEE* 77:2, 1989.
- [17] T. Smith, S. Niekum, D. Thompson, and D. Wettergreen. "Concepts for Science Autonomy During Robotic Traverse and Survey," *Proceedings of IEEE Aerospace*, Big Sky Montana, 2005.
- [18] L. Wasserman, "All of Statistics: A Concise Course in Statistical Inference," Springer, 2004.
- [19] D. Wettergreen, N. Cabrol, V. Baskaran, F. Calderon, S. Heys, D. Jonak, R.A. Luders, D. Pane, T. Smith, J. Teza, P. Tompkins, D. Villa, C. Williams, and M. D. Wagner. "Second Experiments in the Robotic Investigation of Life in the Atacama Desert of Chile," *Proceedings of iSAIRAS*, 2005.
- [20] C. Williams and C. E. Rasmussen. "Gaussian Processes for Regression," *Advances in Neural Information Processing Systems*, 1996.
- [21] M. Varma and A. Zisserman, "A Statistical Approach to Texture Classification from Single Images," *International Journal of Computer Vision: Special Issue on Texture Analysis and Synthesis*, 62:1-2, 2005.
- [22] Y. Yamaguchi and A.B. Kahle and H. Tsu and T. Kawakami and M. Pniel. "Overview of Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER)," *IEEE Transactions on Geoscience and Remote Sensing*, 36:4, 1998.